

USE OF PREDICTIVE PERFORMANCE MODELING DURING LARGE-SCALE SYSTEM INSTALLATION

DARREN J. KERBYSON, ADOLFY HOISIE and HARVEY J. WASSERMAN

*Performance and Architectures Laboratory,
Modeling, Algorithms and Informatics Group, CCS-3
Los Alamos National Laboratory, Los Alamos, NM 87545, USA*

ABSTRACT

In this paper we describe an important use of predictive application performance modeling - the validation of measured performance during a new large-scale system installation. Using a previously-developed and validated performance model for SAGE, a multidimensional, 3D, multi-material hydrodynamics code with adaptive mesh refinement, we were able to help guide the stabilization of the first phase of the Los Alamos ASCI Q supercomputer. We review the salient features of an analytical model for this code that has been applied to predict its performance on a large class of Tera-scale parallel systems. We describe the methodology applied during system installation and upgrades to establish a baseline for the achievable "real" performance of the system. We also show the effect on overall application performance of certain key subsystems such as PCI bus speed and multi-rail networks. We show that utilization of predictive performance models is also a powerful system debugging tool.

Keywords: Performance Modeling; Terascale Systems; Performance Validation; High Performance Computing.

1. Introduction

We have previously reported the development and validation of an analytical model that captures the performance and scaling characteristics of an important ASCI (Accelerated Strategic Computing Initiative) application [1]. We have also described one interesting use of the model to predict the effect on runtime of a key algorithmic change to the application enabling a different parallel decomposition method.

In this paper we report another interesting use of this same model. Los Alamos National Laboratory (LANL) is currently involved in the installation of a Tera-scale computing system called ASCI Q that comprises a large and growing number of compute servers with an interconnect fabric composed of federated switches. ASCI Q will have a peak performance of tens of Teraflops when fully installed during 2003. The installation of a system with such a large number of components is subject to a variety of both hardware- and software-related issues that effectively result in a "stabilization period" during which the systems performance may be sub-par. The question is: how does one identify sub-par performance in a large-scale parallel system, especially one that is larger than any previously available for testing. Performance observations made on a newly-installed system do not

necessarily represent just the cost of processing the workload but often include temporary idiosyncrasies of the hardware and system software, i.e., bugs, faulty or poorly configured hardware components, and so on.

We report here our experiences using a performance model to validate the measured performance during system integration of (part of) ASCI Q. Several sets of measurements of the application performance were made on the system during installation over a period of months. Only after several iterations of hardware refinements and software fixes did the performance of the system achieve the performance predicted by the model. The model did not necessarily reveal precisely what hardware and/or software refinements were needed; however, it was ultimately the only way to determine that no such further refinements were necessary. During this process the model and corresponding system measurements exposed several important performance characteristics associated with the system, such as the effect of PCI bus speed and the effect of multi-rail networks on the overall application performance.

2. The Alpha-Server ES45 Supercomputing System

The system considered here consists of 512 AlphaServer ES45 nodes. Each node contained four 1-GHz Alpha EV68 processors internally connected using two 4-GB/s memory buses to 16 GB of main memory. Each processor has an 8-MB unified level-2 cache, and a 64-KB L1 data cache. The Alpha processor has a peak performance of 2 floating point operations per cycle. Thus this first phase of the Q machine had a peak performance of 4 Tflops. The processors will be upgraded to 1.25-GHz later in the installation process.

Nodes are interconnected using the Quadrics QsNet high-performance network. This network boasts high-performance communication with a typical MPI latency of $5\mu s$ and a peak throughput of 340 MB/s in one direction (detailed measured performance data are discussed in Section 4). The Quadrics network contains two components - the Elan network interface card (NIC), and the Elite switch. The Elan/Elite components are used to construct a quaternary fat-tree topology - an example is shown in Figure 1. A quaternary fat-tree of dimension n is composed of 4^n processing nodes and $n \cdot 4^{n-1}$ switches interconnected as a delta network. Each Elite switch contains an internal 16x8 full crossbar. A detailed description of the Quadrics network can be found in [3].

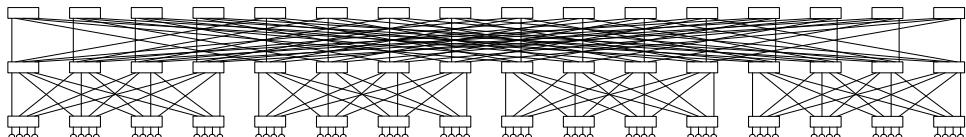


Fig. 1. Network topology for a dimension 3 quaternary fat-tree network with 64 nodes.

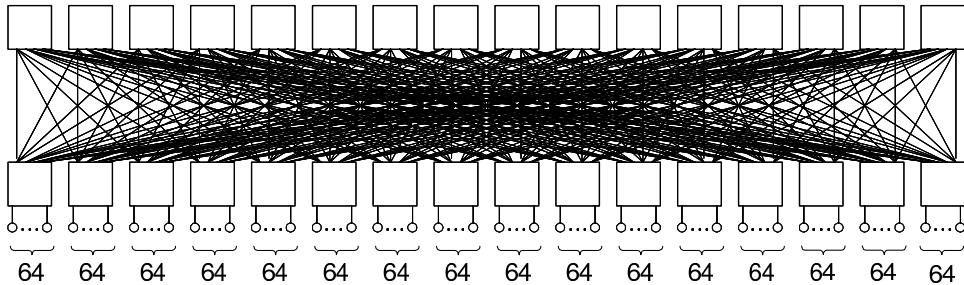


Fig. 2. Interconnection of a federated Quadrics network for a dimension 5 fat-tree.

In order to implement a fat-tree network a single NIC is used per node in addition to a number of Elite switch boxes. The Elite switches are packaged in 128-way boxes which by themselves implement a dimension 3 fat-tree. A 1024-node system, as shown in Figure 2, requires two rows of 16 switch boxes.

Using multiple independent networks, also known as “rails” is an emerging technique to overcome bandwidth limitations and to enhance fault tolerance [5]. The system being installed at Los Alamos contains two rails, i.e. two Elan cards on separate PCI interfaces per node, and two complete sets of Elite switches.

3. The Application and the Model

The application used to analyze the performance of the ES45 cluster is SAGE (SAIC’s Adaptive Grid Eulerian hydrocode). It is a multidimensional (1D, 2D, and 3D), multimaterial, Eulerian hydrodynamics code with adaptive mesh refinement (AMR) consisting of 100,000+ lines of Fortran 90 code using MPI for inter-processor communications. It comes from the LANL Crestone project, whose goal is the investigation of continuous adaptive Eulerian techniques to stockpile stewardship problems. SAGE has also been applied to a variety of problems in many areas of science and engineering including: water shock, energy coupling, cratering and ground shock, stemming and containment, early time front end design, explosively generated air blast, and hydrodynamic instability problems [4]. SAGE represents a large class of production ASCI applications at Los Alamos that routinely run on 1,000s of processors for months at a time.

A detailed description of SAGE, the adaptive mesh processing, and the characteristics of its parallel scaling were described previously in which we developed and validated the performance model [1]. Table 1 gives a summary of the validation results in terms of average and maximum prediction errors across all processor configurations measured. It can be seen that the model is highly accurate with an average prediction error of 5% and maximum of 11% being typical across all machines. These data were not reported previously.

Table 1. SAGE performance model validation results.

System	Configurations tested	Processors tested (maximum)	Error (maximum)	Error (average)
ASCI Blue (SGI O2K)	13	5040	12.6	4.4
ASCI Red (Intel Tflops)	13	3072	10.5	5.4
ASCI White (IBM SP3)	19	4096	11.1	5.1
Compaq AlphaServer ES40	10	464	11.6	4.7
Cray T3E	17	1450	11.9	4.1

4. Use of the SAGE model to validate system performance

The model is parametric in terms of certain basic system-related features such as the sequential processing time and the communication network performance; these had to be obtained via measurements on a small system and are listed in Table 2. The SAGE model is based on weak scaling in which the global problem size grows proportionally with the number of processors. The subgrid size remains constant at approximately 13,500 cells per subgrid.

Table 2. Measured ES45 Performance Parameters for the SAGE Model.

Parameter	33MHz PCI bus	66MHz PCI bus
$T_{comp}(E)(s)$	0.38	0.38
$L_c(S)(\mu s)$	$\begin{cases} 9.0 & S < 64bytes \\ 9.70 & 64 \leq S \leq 512 \\ 17.4 & S > 512 \end{cases}$	$\begin{cases} 6.11 & S < 64bytes \\ 6.44 & 64 \leq S \leq 512 \\ 13.8 & S > 512 \end{cases}$
$1/B_c(S)(ns)$	$\begin{cases} 0.0 & S < 64bytes \\ 17.8 & 64 \leq S \leq 512 \\ 12.8 & S > 512 \end{cases}$	$\begin{cases} 0.0 & S < 64bytes \\ 12.2 & 64 \leq S \leq 512 \\ 8.30 & S > 512 \end{cases}$
$T_{mem}(P)(\mu s)$	$\begin{cases} 1.8 & P = 2 \\ 4.8 & P > 2 \end{cases}$	$\begin{cases} 1.8 & P = 2 \\ 4.8 & P > 2 \end{cases}$

The model is based on a static analysis of the code and includes a summation of the time to process a single cell (multiplied by the number of cells per processor) and the communication of boundary data between processors. The amount of overlap between messaging and computation is negligible. The model is parameterized in terms of dynamic aspects of the code - i.e. those features which are not known through a static analysis. These include: the number of cells, the number of newly adapted cells, and also any load-balancing across processors. These "histories" need to be known on a cycle by cycle basis if the model prediction is to be accurate.

The installation process required that the model predict performance for the first phase of ASCI Q with 2 different PCI bus speeds (initially 33MHz and later 66MHz). The PCI bus speed determines the available bandwidth between the Quadrics NIC and processor memory and has a significant impact on performance. In addition, when two NICs are present within the node (in a 2-rail system), the asymptotic bandwidth increases by approximately 180% if simultaneous messages can take advantage of the two rails [3]. Individual messages are not striped across rails.

4.1. Expected Performance

The performance model was used to provide the expected performance of SAGE on the ES45 system with a 33-MHz and 66-MHz PCI bus using one and two Quadrics rails. These predictions are shown in Figure 3.

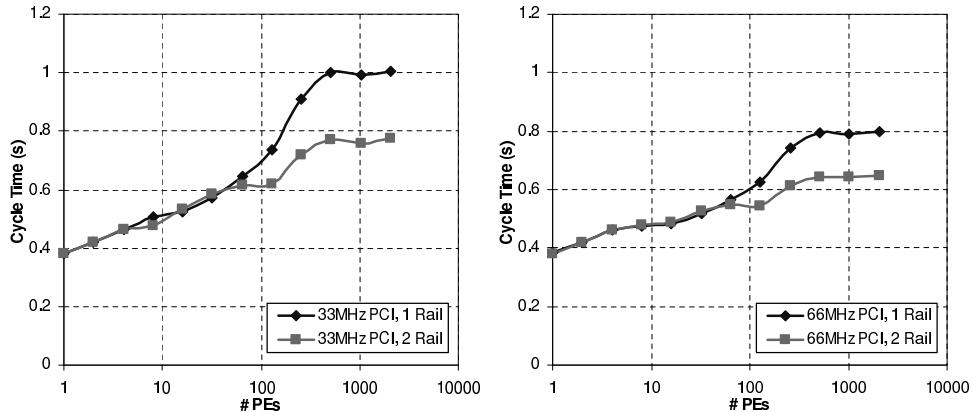


Fig. 3. Performance predictions of SAGE on an ES45 system with QsNet with a) 33-MHz PCI bus, and b) 66-MHz PCI bus.

We note the following observations: 1) since the runs of SAGE were performed for weak scaling, the time should ideally be constant across all processor configurations; 2) the predicted performance is better when using 2 rails than that when using 1 rail. This occurs after a certain point (48 processors) - the point at which a domain is mapped to more than one processor. At this point each node generates more than one simultaneous out-of-node communication for the gather/scatter operations whereas below this point it was only one; 3) the model predicts that the two-fold improvement in PCI bus speed results in only a 20% performance improvement in the code overall; 4) the SAGE cycle time is predicted to plateau above 512 processors this is the point at which all gather/scatter communications are out-of-node.

4.2. Measured Performance

Table 3 summarizes the test conditions on each test date and the refinements made from one test to another. The performance of SAGE was measured at several points after the installation of the machine had taken place: as soon as the machine was up and running (Sept. 9th), after a O/S upgrade from EFT 1.0 to EFT 2.0 and faulty hardware was replaced (Sept. 24th), and after an O/S patch (Oct. 24th). The upgrades that were most significant in terms of performance included bug fixes to the Quadrics RMS resource scheduling software and O/S patches that affected the

priority of a process that determined the allocation of the two rails. Interestingly, this affected both 1- and 2-rail performance. These three sets of measurements, which are based on the 33-MHz PCI bus, are compared with the model in Figure 4 (runtime vs. configuration in Figure 4a and percent error vs. configuration in Figure 4b). The corresponding model prediction and measurements based on the 66-MHz PCI bus are shown in Figure 5. Initially (Jan 4th), not all nodes ran at 66 MHz. By Feb 2nd this had been resolved; however, not all nodes were available for testing. The Quadrics QsNet requires contiguous nodes in order to use its hardware-based collective operations. When nodes are configured out then a software component in the collectives is required which reduces overall performance. By April 20th all nodes were configured in and SAGE achieved the performance predicted by the model at all configurations except for 512 nodes.

Table 3. Summary of Test Conditions.

Date	OS version	No. nodes in system	Performance issues
9-Sept-01	EFT 1.0	128	Some faulty nodes and communication links resulted in poor communication performance, especially on 2 rails
24-Sept-01	EFT 2.0	128	Faulty hardware replaced but still poor 2-rail performance
24-Oct-01	EFT 2.1	128	2-Rail OS patch improved Quadrics Performance
04-Jan-02	EFT 3.0	512	PCI bus upgraded to 66 MHz; SAGE performance at pre-Oct-24 (not all nodes successfully ran at 66 MHz)
02-Feb-02	EFT 3.0	512	All nodes at 66 MHz but some configured out causing lower performance in collective communications
20-April-02	EFT 4.0	512	All nodes configured in, collective performance improved

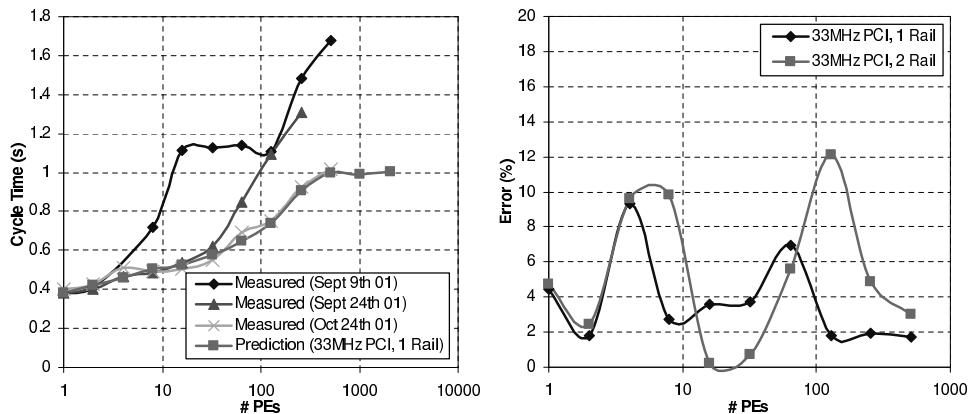


Fig. 4. Measured performance of SAGE (33MHz PCI bus) compared with model predictions. a) Measurement history and model predictions using a single rail, and b) error between final measurements and model when using either 1 or 2 rails.

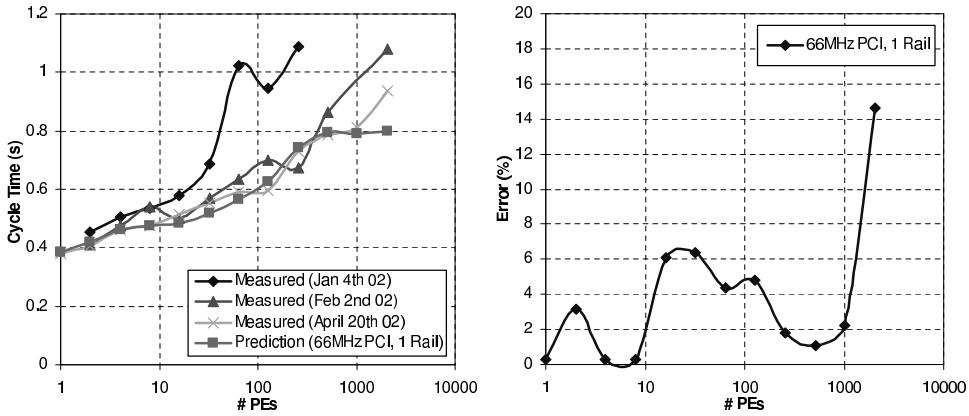


Fig. 5. Measured performance of SAGE (66MHz PCI bus) compared with model predictions. a) Measurement history and model predictions using a single rail, and b) error between final measurements and model on 1 rail.

Figures 4a and 5a show that only after all the upgrades and system debugging had taken place that the measurements matched the expected performance. When the measured data matched the modeled data there was some confidence in the machine performing well. Without the model, it would have been difficult to know conclusively when to stop debugging, or more importantly when not to. When differences did occur between the model and measurements, microkernel benchmarks were run on the computational nodes and the communication network to help identify the source of the problem.

5. Summary

Our teams research over the last few years has focused on the development of analytical performance models for the ASCI workload. It has been said that modeling and predicting the performance of large-scale applications on HPC systems, is one of the great, unsolved challenges for computer science [6]. Clearly, ASCI has a critical need for information on how best to map a given application to a given architecture, and performance modeling is the only means by which such information can be obtained quantitatively. Our approach has been successfully demonstrated for the 100,000-line+ adaptive mesh code reported here, for structured [2], and unstructured mesh transport codes [7], and for a Monte-Carlo particle transport code [8].

The work reported in this paper represents a small but important step in applying our performance models in a very practical way. We expect that ASCI platforms and software will be performance engineered, and that models will provide the means for this. The models can play a role throughout a systems lifecycle:

starting at design when no hardware is available for measurement, in procurement for the comparison of systems, through to implementation / installation, and to examine the effects of updating a system over time. At each point the performance model provides an expectation of the achievable performance with a high level of fidelity. The SAGE performance model has been used for procurement purposes but company-sensitive information precludes disclosure of this in the literature. We can report, as here, how the model becomes the tool for assessing machine performance. Implementation milestone tests related to ASCI Q contractual obligations will be based partially on comparison of observed data with predictions from the SAGE model, in a manner similar to the process described in this paper.

When installing a new system, refinements to both the software system, and hardware components, are often necessary before the machine operates at the expected level of performance. The performance model for SAGE has been shown to be of great use in this process. The model has effectively provided the performance and scalability baseline for the system performance on a realistic workload. Initial system testing showed that its performance was almost 50% less than expected. After several system refinements and upgrades over a number of months, the achieved performance matched exactly the expectation provided by the model. Thus, performance models can be used to validate system performance.

Acknowledgements

Los Alamos National Laboratory is operated by the University of California for the National Nuclear Security Administration for the US Department of Energy.

References

- [1] D.J. Kerbyson, H.J. Alme, A. Hoisie, F. Petrini, H.J. Wasserman and M.L. Gittings, Predictive Performance and Scalability Modeling of a Large-scale Application, in *Proc. SC2001*, Denver, CO, 2001.
- [2] A. Hoisie, O. Lubbeck and H.J. Wasserman, Performance and Scalability Analysis of Teraflop-Scale Parallel Architectures Using Multidimensional Wavefront Applications, *Int. J. of High Performance Computing Applications*, **14** (2000) 330–346.
- [3] F. Petrini, W.C. Feng, A. Hoisie, S. Coll and E. Frachtenberg, The Quadrics Network, *IEEE Micro*, **22(1)** (2002) 46–57.
- [4] R. Weaver, Major 3-D Parallel Simulations, BITS - Computing and communication news, Los Alamos National Laboratory, June/July, (1999) 9–11.
- [5] S. Coll, E. Frachtenberg, F. Petrini, A. Hoisie and L. Gurvits, Using Multirail Networks in High-Performance Clusters, in *Proc. of Cluster2001*, Newport Beach, CA, 2001.
- [6] See <http://perc.nersc.gov/main.htm>.
- [7] D.J. Kerbyson, A. Hoisie, and S.D. Pautz, Performance Modeling of Deterministic Transport Computations, in *Performance Analysis and Grid Computing*, (Kluwer, 2003).
- [8] M. Mathis, D.J. Kerbyson, A. Hoisie and H.J. Wasserman, Performance Modeling of MCNP on Large-Scale Systems, in *Proc. Los Alamos Computer Science Institute Symposium*, Santa Fe, NM, Oct. 2002.